# On Technological Change in Crop Yields

Tor Tolhurst

M.Sc. Student

Food, Agricultural and Resource Economics

University of Guelph

ttolhurs@uoguelph.ca

Alan P. Ker

Chair and Professor

Food, Agricultural and Resource Economics

University of Guelph

aker@uoguelph.ca

## Abstract

Technological change in plant research rarely shifts the entire yield distribution upwards as assumed in the agricultural economics literature. Rather, technologies have been targeted at a specific subpopulation of the yield distribution—for example, drought resistant seeds or so-called *racehorse* seeds—therefore, it is unlikely technological advancements are equal across subpopulations. In this manuscript we introduce a mixture model of crop yields with an embedded trend function in the component means, which allows different rates of technological change in each mixture or subpopulation. By doing so, we can test some interesting hypotheses that have been previously untestable. While previous literature assumes an equivalent rate of technological change across subpopulations we reject the null in 84.0%, 82.3%, and 64.0% of the counties for corn, soybean, and wheat respectively. Conversely, with respect to stable subpopulations through time (i.e. climate change) we reject in only 12.0%, 5.4%, and 4.6% of the counties for corn, soybean, and wheat respectively. These results have implications for modelling yields, directing funds regarding plant science research, and explaining the prevalence of heteroscedasticity in yield data.

## Introduction

Crop yields are agriculture's principle unit of productivity measurement. Since the 1940s agricultural production has experienced marked advances in technology causing yields of staple crops to increase throughout the world. Figure 1 illustrates these remarkable increases for Iowa with state-average yields for three staples crops—corn, soybean and wheat—mapped over time. Even the lowest rate of technological advancement, average soybean yields, have more than doubled. On the other end of the spectrum, average corn yields have more than tripled in both areas. Elsewhere crop yields have risen at similarly remarkable rates.[1]

The rate of technological change has been exclusively measured at the mean despite evidence from agronomists (i.e. Khiari *et al.*, 2001) suggesting crop yields may have distinct subpopulations. Technological developments in plant research rarely shift the entire yield distribution upwards but instead target a subpopulation of the yield distribution—for example, drought-resistant seeds versus so-called *racehorse* seeds (which attempt to maximize yields under optimal conditions). In particular, it is common to consider two subpopulations: a *regular* year subpopulation and a *poor* year subpopulation. The existence of two subpopulations would align with empirical evidence from the agricultural economics literature, which suggests the majority of yield distributions are negatively skewed (see in particular Day, 1965; Gallagher, 1987; and Goodwin and Ker, 1998).[2]

We model crop yields as having two subpopulations or components -- a *regular* year component and a *poor* year component -- using a mixture of two normals: $(1 - \lambda)N(\mu_1, \sigma_1^2) + \lambda N(\mu_2, \sigma_2^2)$ where $\lambda$ is the mixture weight, $\mu_j$ is the mean of component $j$ and $\sigma_j^2$ is the variance of component $j$.[3] Using mixtures is not entirely new as others have first estimated

---

[1]Despite these increases there are concerns technological advancements may be stagnating. For example, Ray *et al.* (2012) found 29.9% of corn, 24.3% of soybean, and 38.8% of wheat yields around the world are *"not improving"* (p. 5).

[2]This skewness is often explained as the result of the asymmetric impact of weather shocks on yields: ideal growing conditions cannot generate yields above biological limits while catastrophic growing conditions can reduce yields to zero.

[3]While we denote the components *regular* and *poor* given the tendency towards negative skewness in the yield data, if the underlying data is positively skewed the *poor* component will be located on the upper tail of the *regular* component.
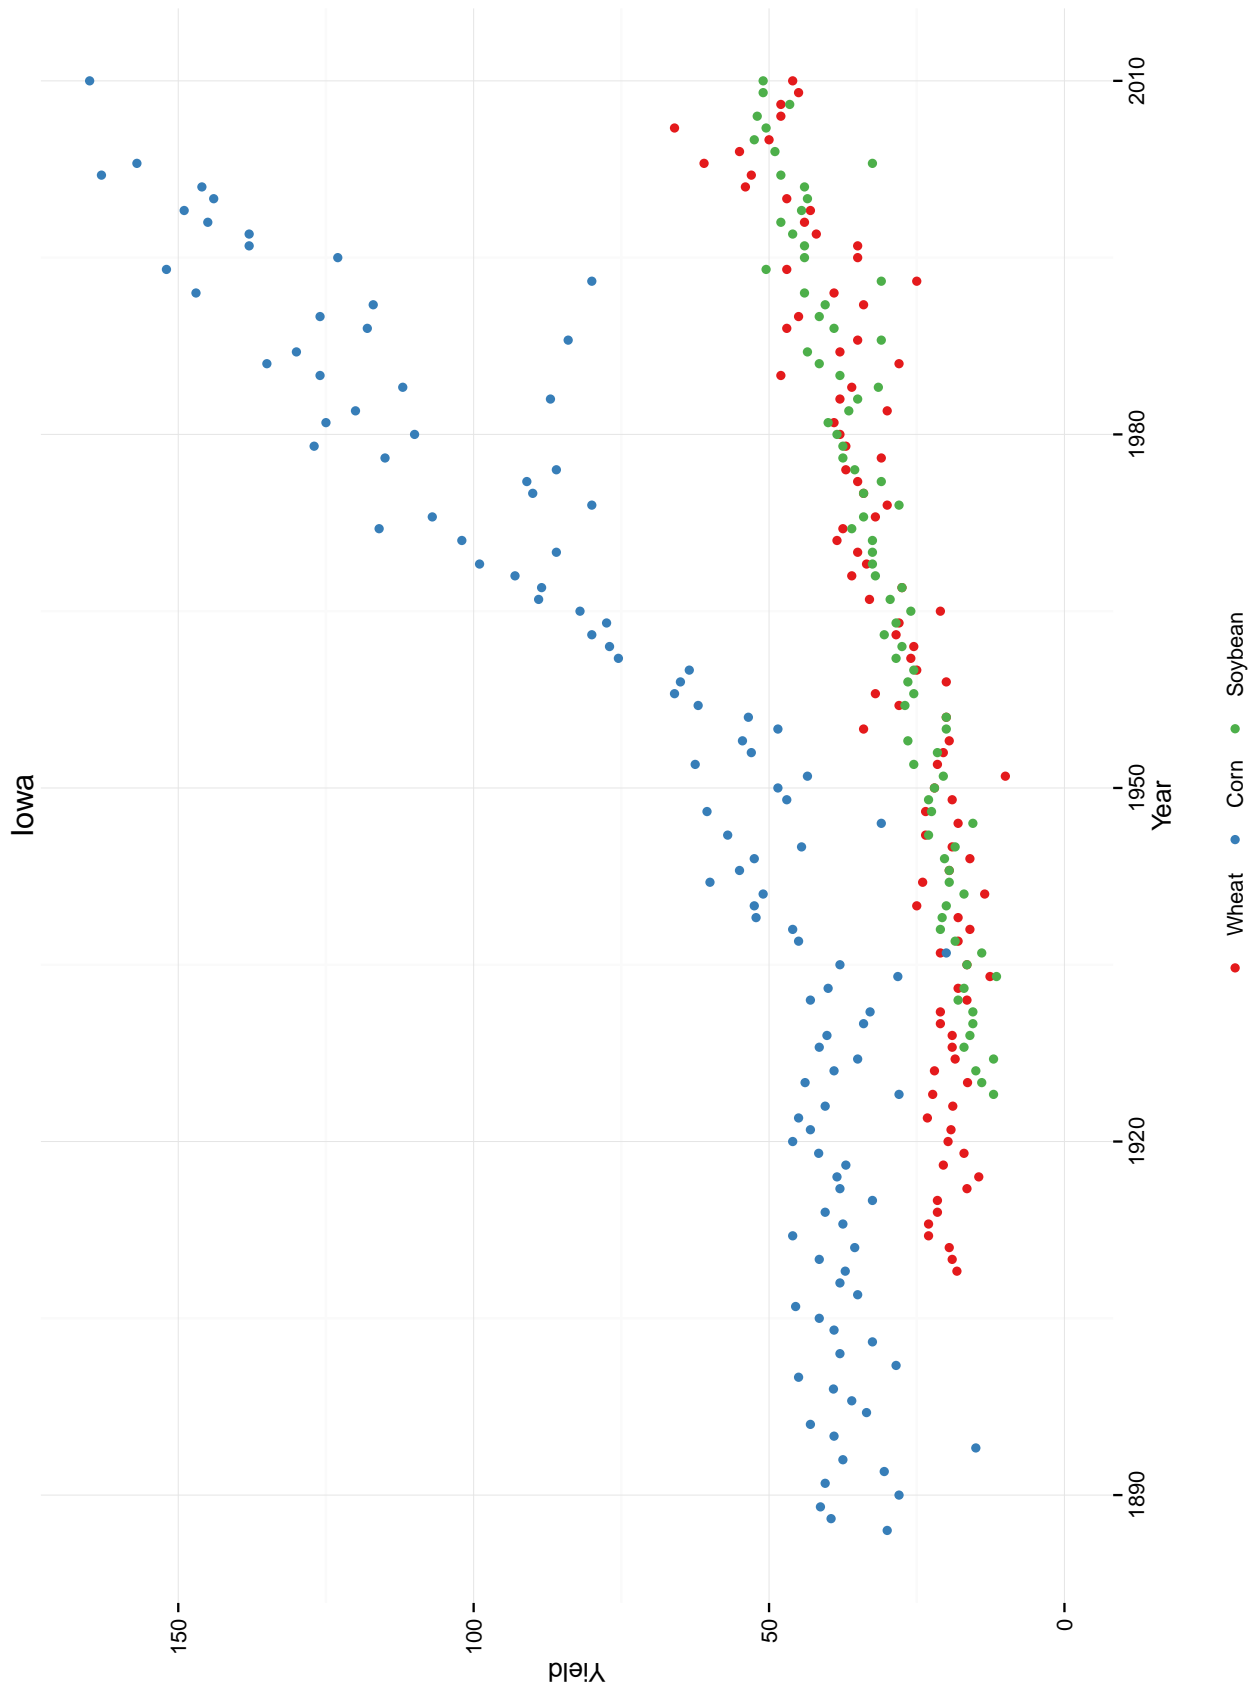
Figure 1: Average Iowa Crop Yields 1887-2010 (National Agricultural Statistics Service).

$y_t = f(t) + \epsilon_t$ and then using $\hat{\epsilon}_t$ estimate a mixture of two normals (e.g. Ker 1996; Goodwin, Roberts and Coble, 2000; Woodard and Sherrick, 2011). However, the mixture model is significantly more flexible than previously employed: the mixture model will allow for unique trend components, or different rates of technological change, within each individual component. That is:

$$y_t \sim (1 - \lambda)N(f(t), \sigma_1^2) + \lambda N(g(t), \sigma_2^2). \tag{1}$$

where $f(t)$ and $g(t)$ are time trend functions.

To illustrate (and provide some intuition) for the proposed model, the estimate for corn yields from Harrison county, Illinois is presented in Figure 2. In this crop-county combination, the rate of technological change in *regular* years appears greater than the rate of technological change in *poor* years. Figure 2b illustrates the estimated yield densities at four different time periods. The shape of the estimated yield densities changes noticeably over time. In 1950 the estimated yield density appears relatively normal, while in 1970 the estimated density displays significant negative skewness. As time increases and the effect of differing rates of technological change become more prevalent, the estimated densities become increasingly bimodal and the overall variance increases (giving rise to the presence of heteroscedasticity).

In contrast, Figure 3 illustrates the estimated trend and mixture densities for the same yield series using the typical method. That is, by first estimating a single temporal process (here a linear trend) and then estimating the mixture from the heteroscedasticity-corrected residuals as in Ker (1996), Goodwin, Roberts and Coble (2000) and Woodard and Sherrick (2011). The estimated densities are, somewhat surprisingly, quite different. Crop insurance premium rates based on the two density estimates at the 75% coverage level are 11.1% for the traditional method and 13.0% for the trend mixture model. At the 90% coverage level, these rates are 31.5% and 28.8%, respectively. Although nominally the difference may appear small, the premium rates from the trend mixture model are 1.17 and 1.09 times larger than the traditional model, respectively. In the context of crop insurance, such differences in rates can have significant economic consequences for the actuarial soundness of a crop insurance

program.

The proposed model can provide a great deal of insight into the relationship between the rate of technological change under *regular* and *poor* years. The presence of differing trend functions would suggest that technology has increased *regular* year yields differently than *poor* year yields. Also, differing rates of technological change may give rise to question previous analyses that condition yields with respect to time at the mean only because these necessarily assume identical rates of technological change and a stable error distribution. Finally, the hypothesis that rates of technological change are identical under *regular* and *poor* years not only has implications for modeling yields, but could also have implications for rating crop insurance contracts, risk management policy, and directing funds regarding plant science research. This model also enables us to test interesting climate change questions such as whether the probability of a *poor* year is increasing or decreasing over time? This too has implications beyond yield modeling.

The manuscript attempts to make three contributions to the literature. First, we develop a model of crop yields that can accommodate differing rates of technological change in different components or subpopulations. Second, we use this model to empirically test if rates of technological change differ between the components. Third, we empirically test if the probability of the components is stable over time. To the best of our knowledge these contributions are new to the literature.

## Crop Yield Models in the Literature

Most often the approach to estimating conditional yield densities is to: (i) estimate a trend using the yield data; (ii) test the residuals from (i) for heteroscedasticity and adjust if necessary; and (iii) estimate a parametric or nonparametric conditional yield density given (i) and (ii). The choice of density estimation method has received by far the most attention in the literature.

A wide variety of density estimation approaches have been proposed in the literature.

5
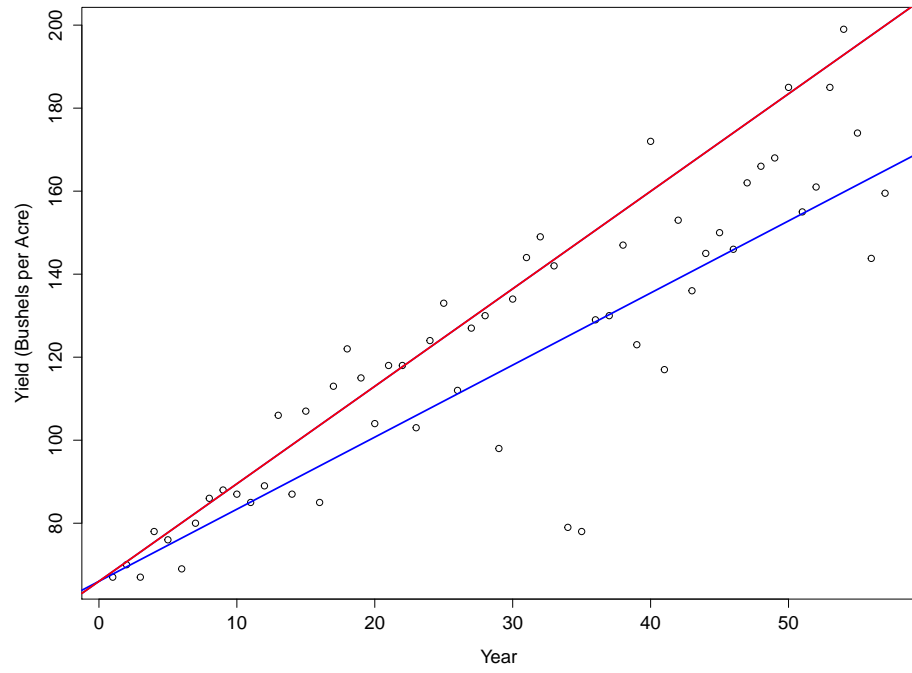
(a) Estimated mixture of two Normals with differing temporal components.
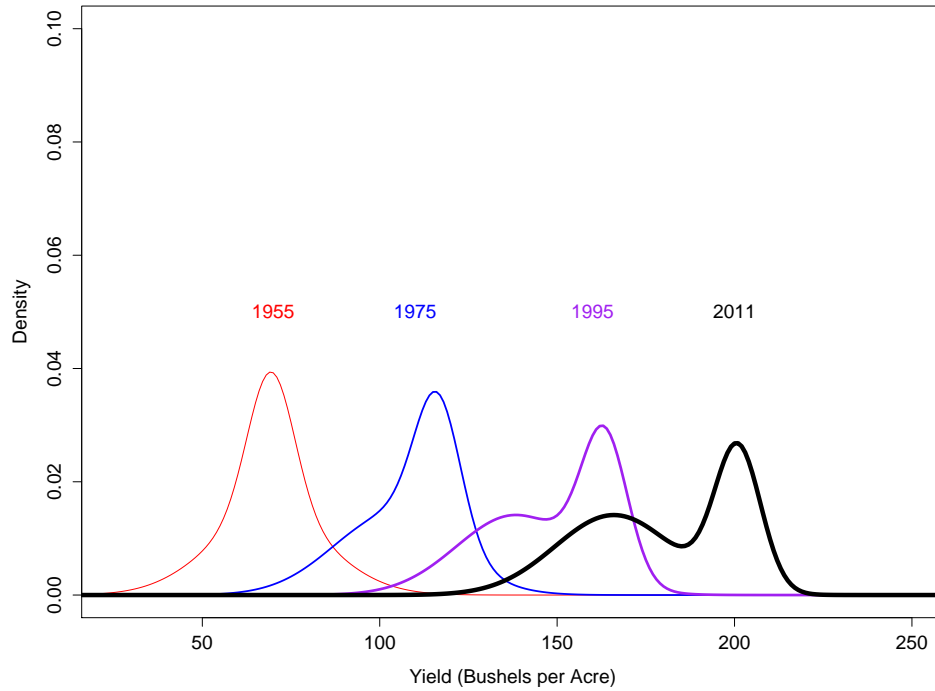


(b) Corresponding conditional yield density estimates over time.

Figure 2: County-level corn yields in Harrison, Illinois with estimated two-component trend.

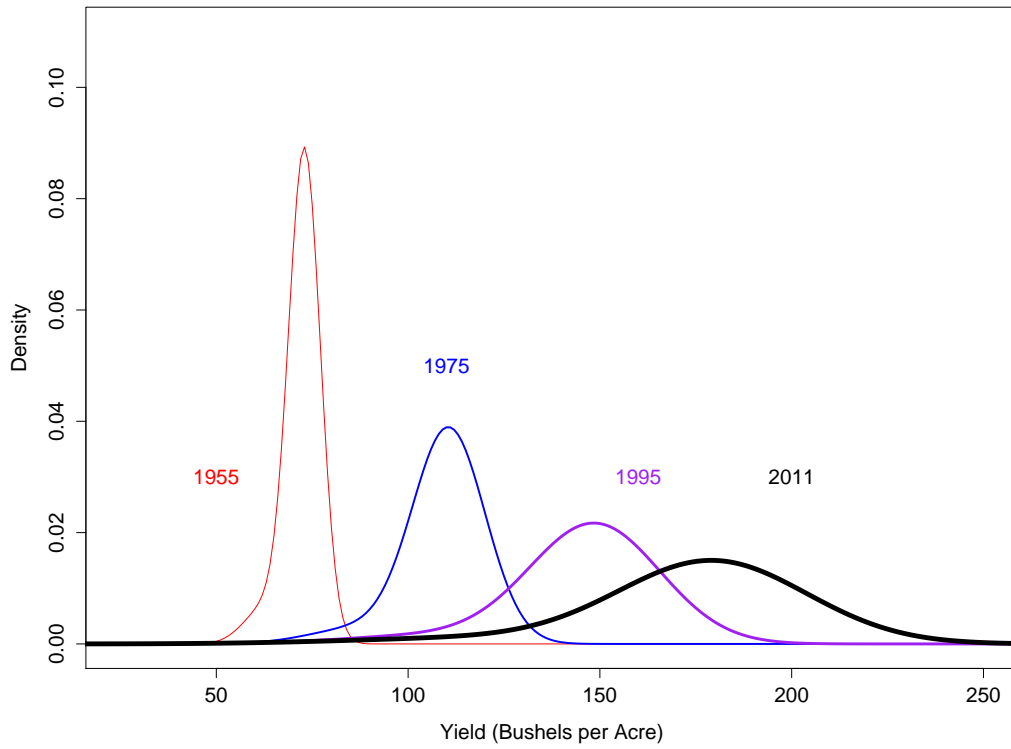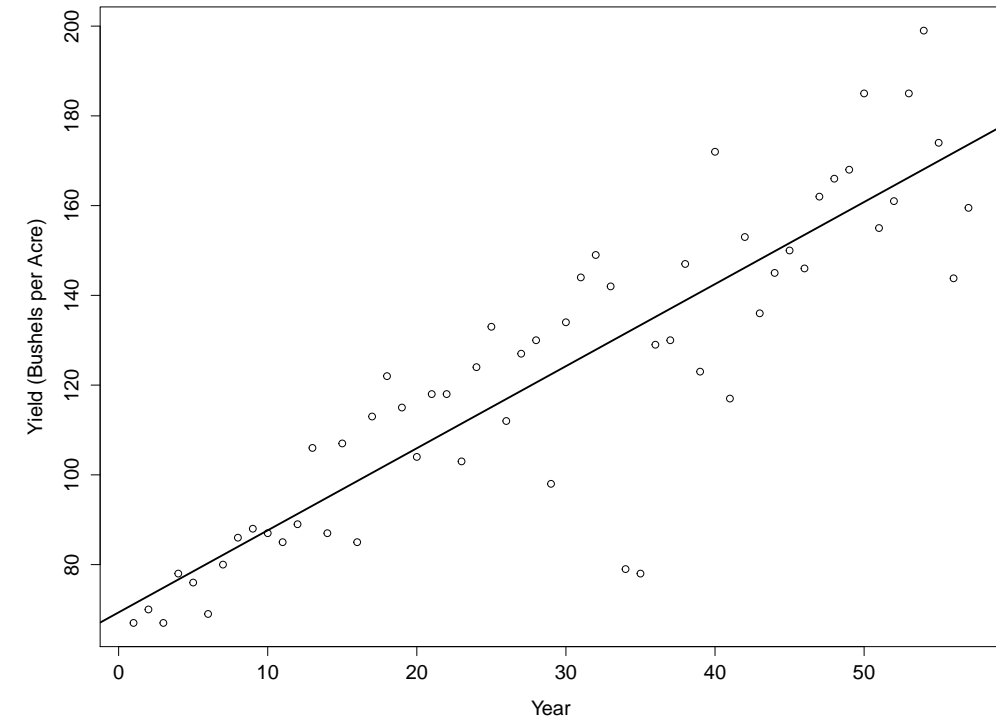Figure 3: County-level corn yields in Harrison, Illinois and estimated densities following the traditional estimation procedure.

In 1958, Botts and Boles first suggested the use of "normal curve theory" to determine crop insurance premium rates. The seminal contribution of Day (1965) argued crop yield densities displayed non-normal attributes such as significant skewness. In response, Gallagher (1987) suggested the gamma distribution while Nelson and Preckel (1989) suggested the beta distribution. Goodwin and Ker (1998) proposed nonparametric kernel density methods while Just and Weninger (1999) argued deviations from normality were the result of inconsistencies in methods and data. A semi-parametric approach was forwarded by Ker and Coble (2003). Later parameteric specifications included the logistic (Atwood, Shaik and Watts, 2003) and Weibull distributions (Sherrick *et al.*, 2004). Inverse sine transformation methods were used by Moss and Shonkwiler (1993), Ramirez (1997), Ramirez, Misra and Field (2003), Ramirez and McDonald (2006). The model proposed here is similar in that it uses a mixture of two Normals, which has been used to model yields elsewhere (Ker, 1996; Goodwin, Roberts and Coble, 2000; and Woodard and Sherrick, 2011). More than two mixtures can used; however two are sufficiently flexible to accommodate a variety of distributional structures that are commonly associated with yield data such as symmetry or skewness (both negative and positive), long-tailed and bimodal. Notably, estimated yield densities using a mixture of two Normals are nearly identical to estimates from nonparametric kernel methods.

Trend estimation and heteroscedasticity correction have received far less attention in the literature. Both deterministic and stochastic approaches have been considered in estimating the rate of technology change or trend commonly present in yield data. Deterministic approaches have dominated the literature and include a simple linear trend, two-knot linear spline (Skees and Reed, 1986), and polynomial trend (Just and Weninger, 1999). Stochastic approaches include the Kalman filter (Kaylen and Koroma, 1991) and ARIMA$(p, d, q)$ (Goodwin and Ker, 1998). More recently, Ozaki and Silva (2009) and Claassen and Just (2011) incorporated spatial information into their temporal model. In contrast, we are proposing a model that allows different rates of technological change in different yield subpopulations. With the exception of Harri *et al.* (2011) and Just and Weninger (1999), heteroscedasticity

has received surprisingly little attention in the literature considering the magnitude of its effect on crop insurance rates. Interestingly, different rates of technological change are consistent with the wide-spread prevalence of heteroscedasticity in crop yields: if *regular* year yields increase at a higher rate than *poor* year yields, the dispersion of crop yields would increase through time despite homoscedastic component variances.[4]

## Empirical Approach

Two caveats about crop yield data are necessary. First, it is important to use the lowest level of data aggregation possible because when yield data is aggregated it tends to loose some of the nonnormal characteristics we described earlier (i.e. skewness, bimodality). Due to the limited availability of farm-level yield data with a sufficiently long time series we use county-level data.[5] Second, realized yield data are a function of adopted technologies and not necessarily the set of technologies available to producers. Therefore, our conclusions concern the rate of *adopted* technological change rather than *possible* technological change.[6]

We estimate the model using county-level corn, soybean and wheat yield data. Corn and soybean data are from Illinois, Indiana and Iowa for 1955 - 2011. All three of these states are major producers of both corn and soybeans: in 2011 they accounted for 15.6%, 7.2% and 17.3% of national corn production and 13.7%, 7.8% and 15.4% of national soybean production in the United States, respectively (NASS, 2013). These states were ranked second, fifth and first in total bushels of both corn and soybeans produced in the United States, respectively. For wheat, we use Kansas, Nebraska and Texas for 1968 - 2008. The time series is shorter due to limited availability of county-level data for winter wheat. These three states were also major producers, accounting for 24.2%, 4.3% and 8.6% of national wheat production in 2010,

---

[4]Or potentially *vice versa* with decreasing variance over time.

[5]Extrapolating conclusions from county-level yield data to the farm-level can lead to incorrect conclusions and must be done with caution. See especially Just and Weninger (1999).

[6]For example *poor* year yields may not be increasing at the same rate as *regular* year yields because producers increasingly choose *racehorse* seeds over *workhorse* seeds and seek downside protection through financial mechanisms such as crop insurance. The county-level yields are the *net aggregate* of these choices for the area and don't necessarily speak directly to individual producer choices.

respectively. These states were ranked first, sixth and second rank in national production wheat. All 791 crop-combinations of data for the United States were collected from the U.S. Department of Agriculture's National Agricultural Statistics Service. Any county without a full yield history was excluded.

As noted above, we consider a mixture of two normals where the mean of each normal is not static but rather represents the temporal process of technological change in each component or mixture. That is, we have

$$y_t \sim (1 - \lambda)N(\alpha_P + \beta_P t, \sigma_P^2) + \lambda N(\alpha_R + \beta_R t, \sigma_R^2). \tag{2}$$

The unknown parameters to be estimated are $\lambda$, $\alpha_P$, $\beta_P$, $\sigma_P^2$, $\alpha_R$, $\beta_R$, and $\sigma_R^2$. Mixture models are commonly estimated using the Expectations-Maximization (EM) algorthm because of convergence issues with likelihood methods. We necessarily modified the traditional form of the EM algorithm (Dempster, Laird and Rubin, 1977) to embed the temporal functions. This requried replacing the weighted means estimate for updating the component means in the traditional algorithm with a weighted least squares estimate. The diagonal of the weighting matrix for the weighted least squares estimate is the weighting vector in the traditional EM algorithm.

Mixture models are commonly estimated using the expectation-maximization (EM) algorithm because of convergence issues with likelihood methods. In order to estimate the component trend, we modified the traditional form of the EM algorithm of Dempster, Laird and Rubin (1977) by embedding the trend functions in the component means. For estimation, this required replacing the weighted means estimate (for updating the component means in the traditional algorithm) with a weighted least squares estimate. The diagonal of the weighting matrix for the weighted least squares estimate is the weighting vector in the traditional EM algorithm.

Karlis and Xekalaki (2003), citing Bohning (1999) and McLachlan and Peel (2000), sum-

marize the advantages and disadvantages of the EM algorithm. Among the notable advantages, the EM algorithm "exhibits monotonic convergence" and "leads to estimates within the admissible range if the initial values are within the admissible range" (p. 578). The monotonic convergence implies parameter estimates improve (or do not regress) at each iteration. The main limitation of the algorithm is that it may converge on local maxima, particularly when the log-likelihood function is relatively flat or has multiple peaks. The problem of local maxima can be reduced by choosing multiple starting values. Starting values may be either chosen for the parameters or for the probability that a given realization belongs to a given component. We attempted three different approaches and found identical results in almost all cases.[7] [8]

## Estimation Results

Figure 4 presents the estimated temporal process for representative county-crop combinations. For comparison, Figure 5 presents results selected to illustrate some of the more extreme cases. Consider in particular, Cherokee soybean. The Cherokee soybean case illustrates the atypical case when $\hat{\beta}_R < \hat{\beta}_P$ and the variance of yields appears to be decreasing. Interestingly this case runs counter to the majority of empirical findings, which typically suggest variance is either constant or increasing over time. The atypical cases provide some interesting results; however the vast majority of counties are consistent. In the vast majority of cases the rate of technological change is higher in *regular* year yields than in *poor* year yields.

A more complete picture of the relationship between $\hat{\beta}_R$ and $\hat{\beta}_P$ is provided in Figure 6 which maps $\hat{\beta}_R$ against $\hat{\beta}_P$ for all counties of corn, soybean and wheat respectively. The

---

[7]First we assigned a given yield realization probability zero to the *poor* component if it was greater than one standard deviation below the mean trend and one otherwise. Second we assigned a given yield realization probability zero to the *poor* component if it was below the mean trend and one otherwise. Third we choose starting values for the parameters $\lambda$, $\alpha_P$, $\beta_P$, $\sigma_P^2$, $\alpha_R$, $\beta_R$, and $\sigma_R^2$.

[8]In cases where there where no significant *poor* yields in the first 5-10 years, the trend line for the *poor* component crossed the trend line in the *regular* component—not surprising given least squares. In these cases, we constrained the intercept from the *poor* component to be equal to the *regular* component to prevent the trend lines from crossing.
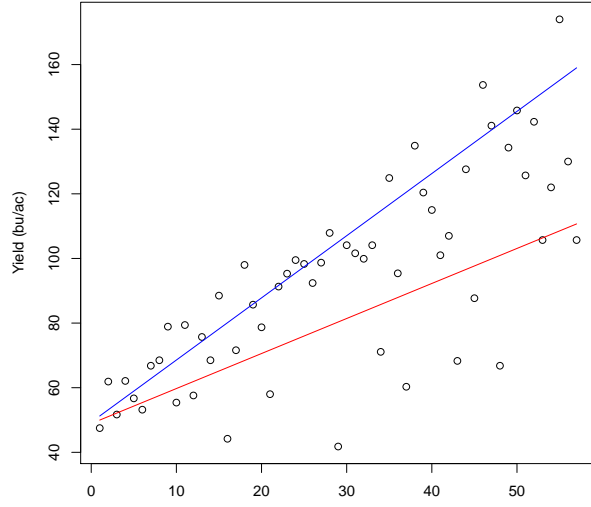
(a) Harrison, Indiana Corn
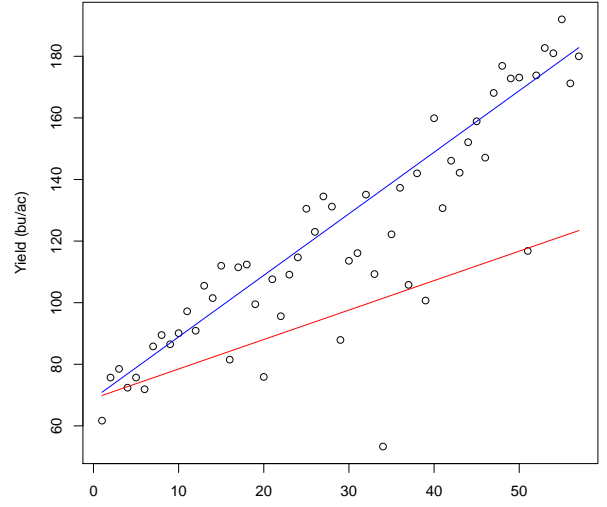
(b) Clinton, Iowa Corn
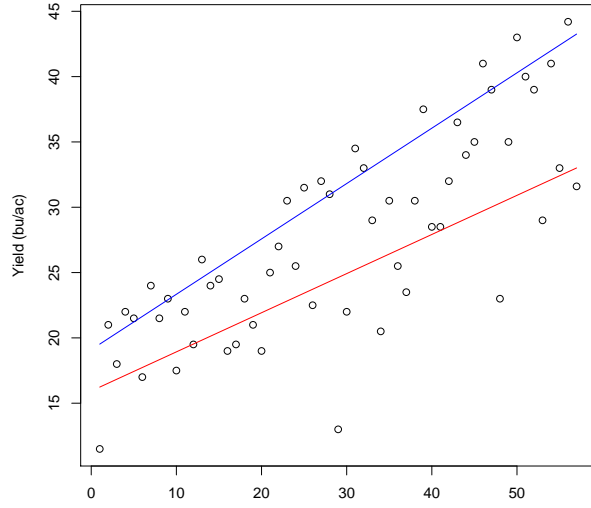
(c) Clay, Illinois Soybean
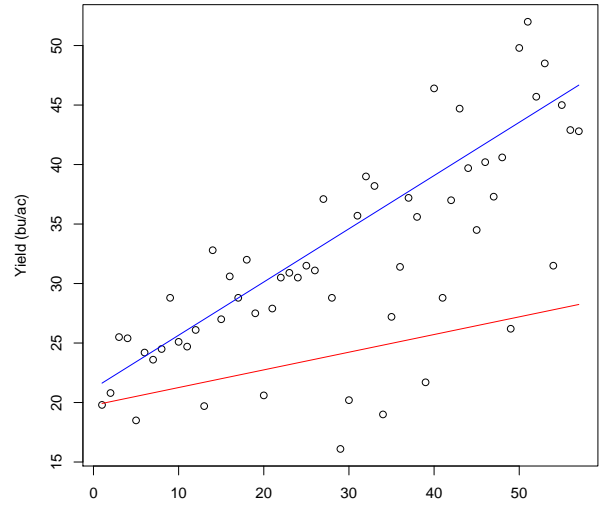
(d) Decatur, Iowa Soybean
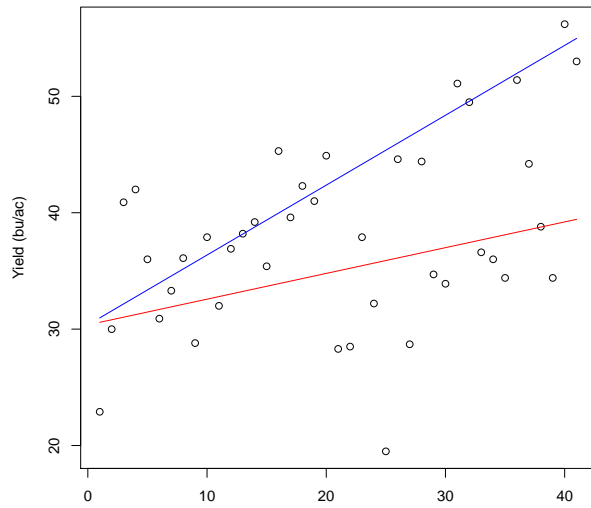
(e) Keith, Nebraska Wheat
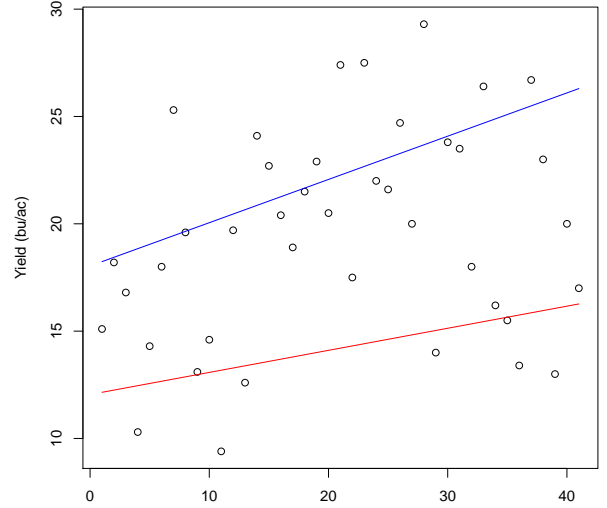
(f) Scurry, Texas Wheat

Figure 4: Representative two-component technological trend estimates.
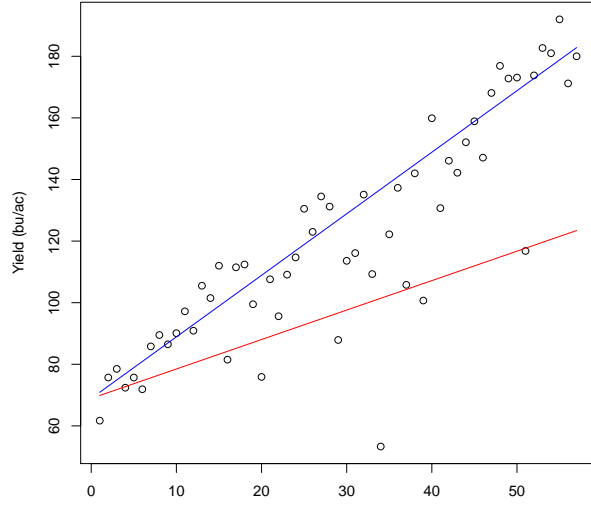
(a) Clinton, Iowa Corn
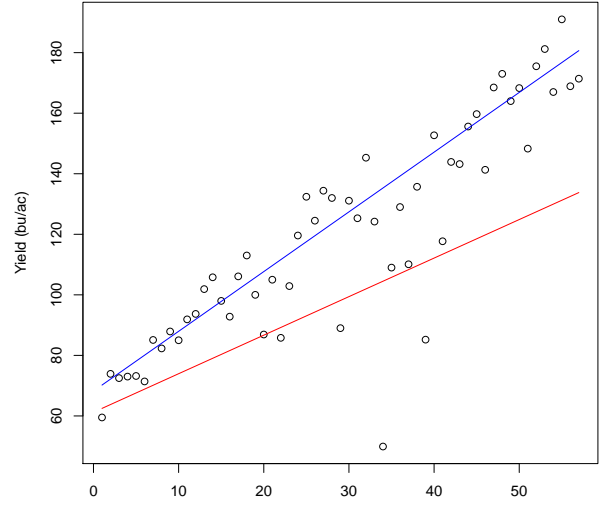
(b) Jones, Iowa Corn
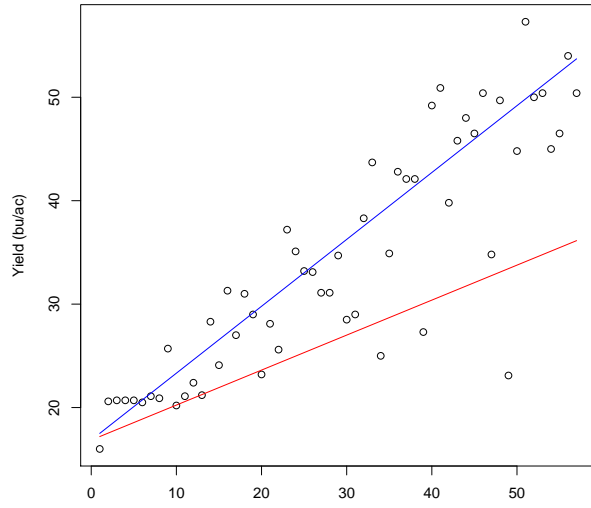
(c) Chickasaw, Iowa Soybean
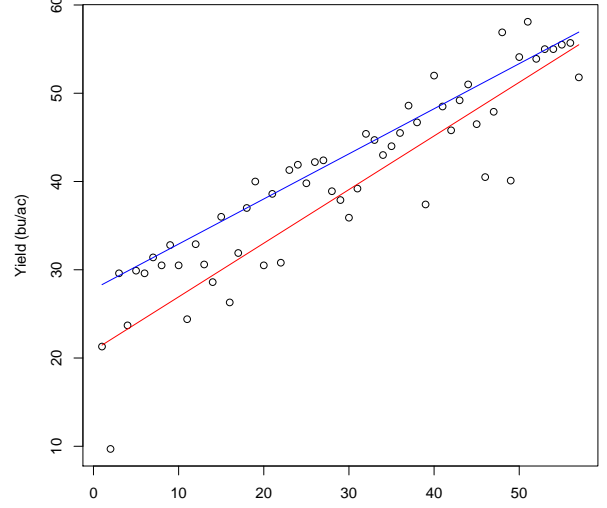
(d) Cherokee, Iowa Soybean
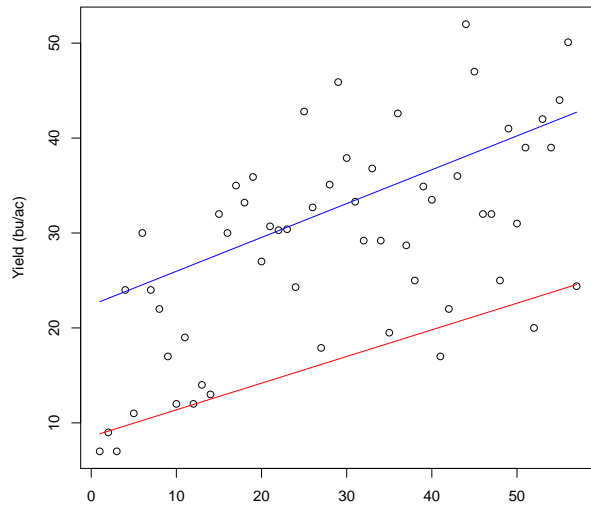
(e) Hodgeman, Kansas Wheat
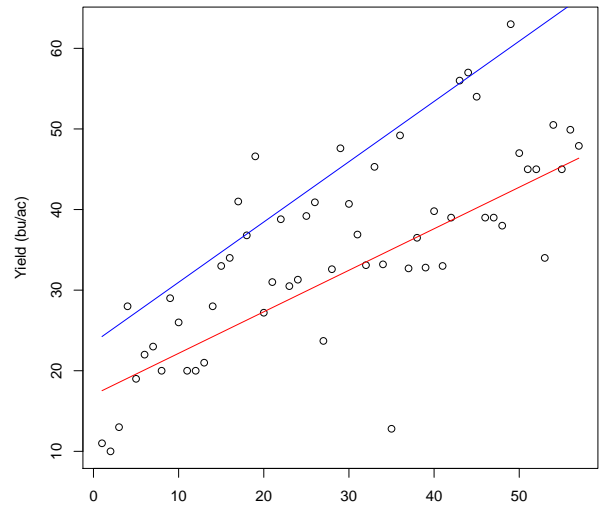
(f) Mitchell, Kansas Wheat

Figure 5: Atypical two-component technological trend estimates.
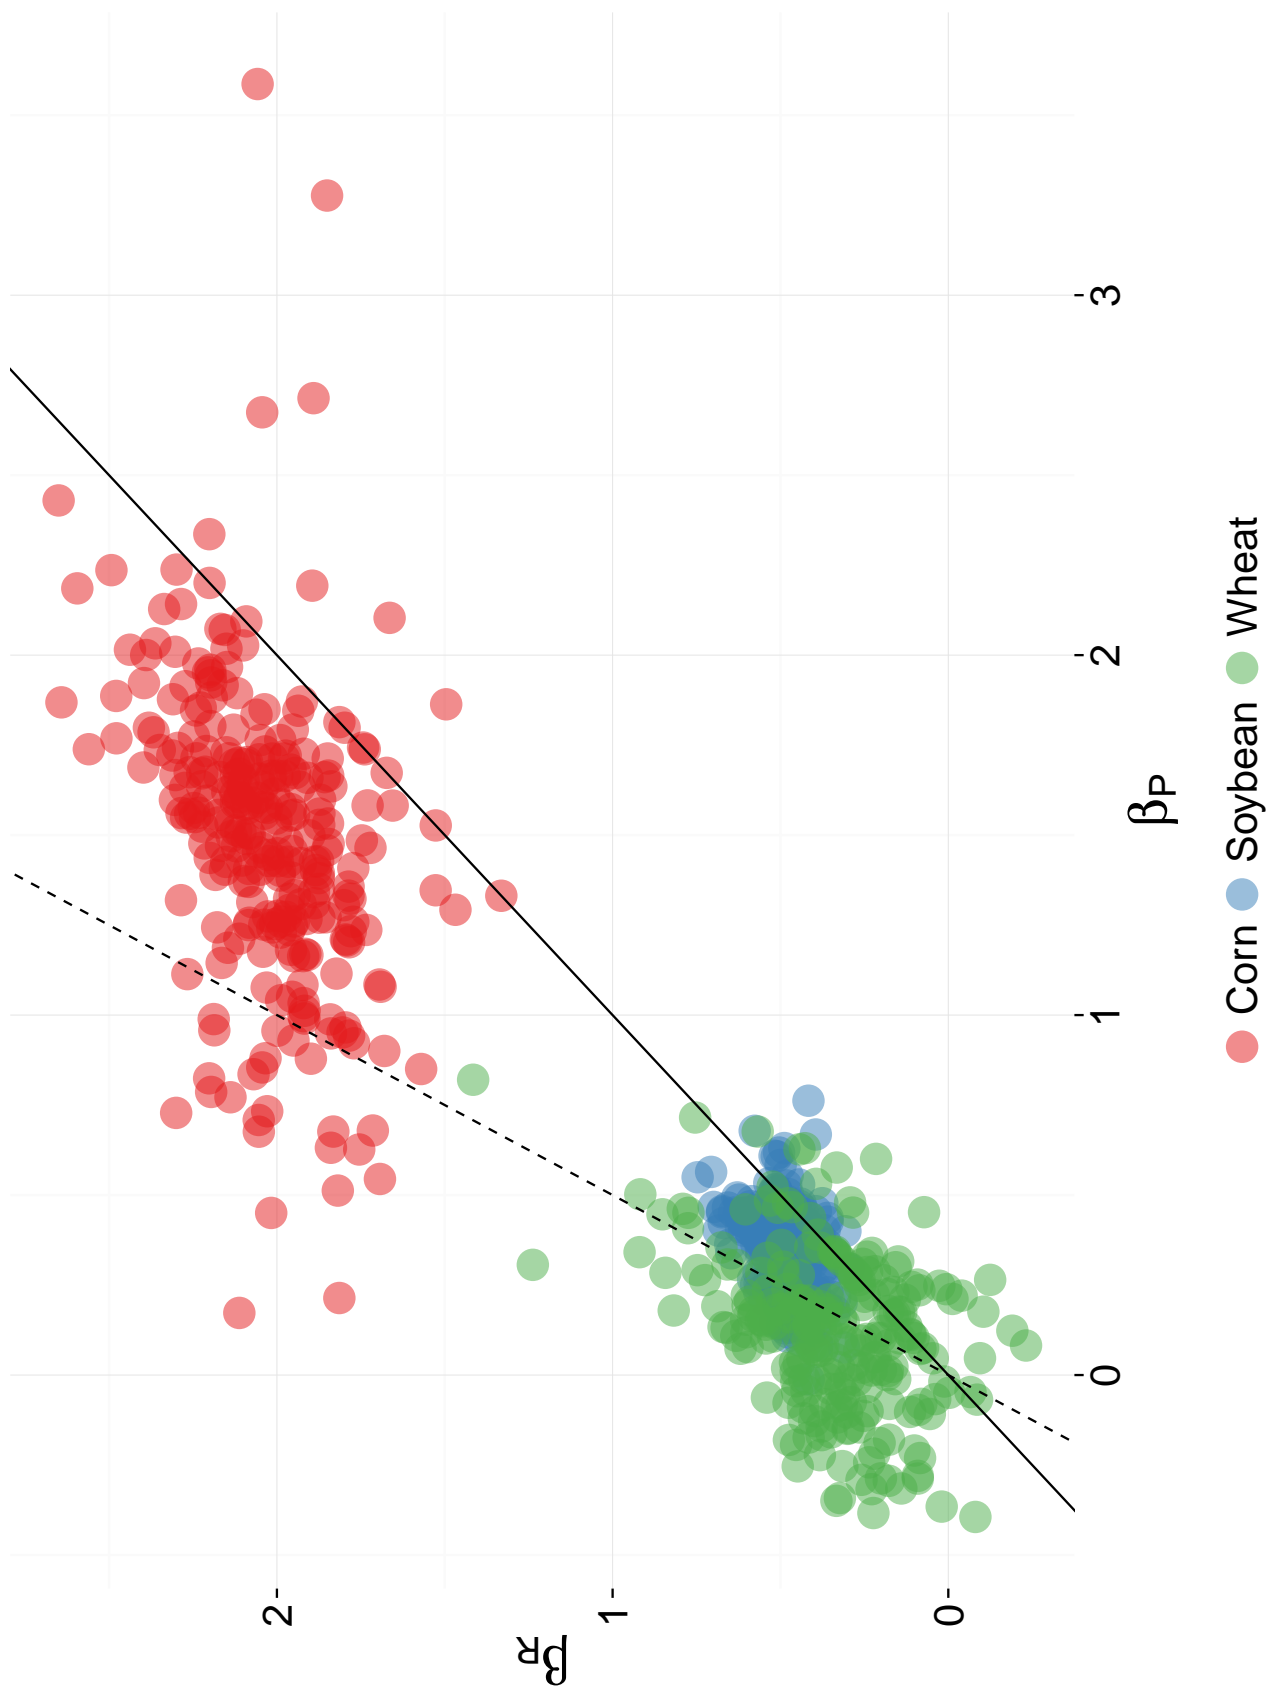
Figure 6: Rate of technological change in *regular* year versus *poor* year yields across all states.

solid line illustrates when the rates of technological change are equivalent in *regular* and *poor* year yields. Notably, this solid line corresponds to the assumption of the current literature. First off, it is readily apparent the rate of technological change in corn—regardless of *regular* or *poor* year—has significantly outpaced soybean and wheat. For corn, $\hat{\beta}_R$ is never below one and the majority of $\hat{\beta}_P$ exceed one. For soybean and wheat, in contrast, even $\hat{\beta}_R$ rarely exceeds one. It is also clear in these illustrations the rate of technological change in *regular* year yields has outpaced the rate in *poor* year yields by a considerable margin in the vast majority of cases for all three crops: the vast majority of points lie above the solid line. Only a small number of cases have $\hat{\beta}_R < \hat{\beta}_P$ and fall below the solid line: 5.3% of corn, 8.7% of soybean and 6.2% of wheat counties. Overall in 88.5% of crop-county combinations the rate of technological change is higher in *regular* year yields than in *poor* year yields. The dashed line illustrates when $\beta_R$ is two times $\beta_P$ and in a number of cases the rate of technological change in *regular* year yields has doubled the rate of change in *poor* years, particularly in soybean and wheat where nominal yields are low. Table 1 reports summary statistics of the ratio in the estimated rate of technological advancement in *regular* years over the rate of technological advancement in *poor* years broken down by crop and regions. Not surprisingly, the mean $\hat{\beta}_R/\hat{\beta}_P$ ratio is above one, reinforcing the difference in the rates of technological change.

Also reported in Table 1 is a more statistically rigorous analysis of the empirical question: has technological advancement effected *regular* year yields and *poor* year yields at the same rate? We evaluate this question with a likelihood ratio test under the following hypothesis:

$$H_o^1 : \beta_P = \beta_R$$
$$H_a^1 : \beta_P \neq \beta_R$$

The results of this likelihood ratio test provide further evidence of a different rate of technological advancement *regular* and *poor* year yields, which is prevalent in the vast majority

Table 1: Hypothesis Test One Results.

| Crop | Summary Statistics for Ratio of $\hat{\beta}_R/\hat{\beta}_P$ Ratio | | | | | Rejection Rate of $\hat{\beta}_R = \hat{\beta}_P$ |
|---|---|---|---|---|---|---|
| | Minimum | Mean | Median | Maximum | Std. Dev. | |
| **Corn** | | | | | | |
| Illinois | 0.56 | 1.35 | 1.34 | 2.52 | 0.23 | 81.4% |
| Indiana | 0.70 | 1.53 | 1.38 | 8.48 | 0.89 | 84.8% |
| Iowa | 0.57 | 1.69 | 1.32 | 12.24 | 1.25 | 85.9% |
| **Soybean** | | | | | | |
| Illinois | -34.20 | 1.12 | 1.38 | 4.68 | 3.68 | 83.5% |
| Indiana | 0.92 | 1.39 | 1.34 | 3.03 | 0.36 | 80.5% |
| Iowa | 0.82 | 1.79 | 1.51 | 4.60 | 0.86 | 82.7% |
| **Winter Wheat** | | | | | | |
| Kansas | -642.27 | -5.59 | 0.82 | 129.75 | 68.49 | 79.6% |
| Nebraska | -26.79 | 1.57 | 1.51 | 15.11 | 5.64 | 56.0% |
| Texas | -25.65 | 1.21 | 1.00 | 43.56 | 6.97 | 53.1% |

Note: Rejection rate is the per cent of counties rejecting the null hypothesis evaluated at the 5% significance level. The extreme values (for example a maximum ratio of 129.75 and a minimum ratio of -642.27 for Kansas wheat) are extreme because $\hat{\beta}_P \to 0$, which inflates the ratio. These values are extremely high when they approach zero from the righthand side and extremely low when they approach zero from the lefthand side. These values are apparent in Figures ?? to ?? near the vertical axis and there is nothing to indicate they are problematic.

of crop-county combinations. In the vast majority of cases the null hypothesis is rejected: 84.0% for corn, 82.3% for soybean and 64.0% for wheat.[9] These results call to question the approach used by Ker (1996), Goodwin, Roberts and Coble, (2000), and Woodard and Sherrick, (2011) and more generally, the time-conditioning of yields exclusively at the mean in crop yield models.

Hypothetically, technological change could also effect the crop yield distribution through the probability of a *regular* year parameter. Ideally, technological change would result in more resilient crops and production practices that would increase the probability of a *regular* year. With this in mind, the second hypothesis test evaluates the empirical question: has the probability of a *regular* year changed over time? To test this hypothesis we regress $\hat{\gamma}_t$ (the probability of membership) against $t$ in order to examine if the time coefficient is significantly different from zero using a $t$-test with robust standard errors. Table 2 presents a summary of the $t$-test results. For the most part, the vast majority of county-crop combinations fail to reject the null hypothesis: 12.0% for corn, 5.4% for soybean and 4.6% for wheat. The most compelling results—which suggest the probability of for the three American states may be decreasing—fail to reject in a very high proportion of the counties. Overall, the results suggest very little evidence the temporal process of technological change has significantly effected the probability of *regular* year.

## Economic Implications for Crop Insurance

The purpose of this section is to provide an example of where the proposed empirical model would have economic implications. To do so, we compare the crop insurance rates estimated from mixture models with both differing rates of technological change, denoted two-trend model, and a single rate of technological change, denoted simple-trend model. To compare rates, the out-of-sample simulation compares two crop insurance rating processes using historical data. The simulation involves two agents: (1) the Risk Management Agency (RMA)

---

[9]Interestingly, corn and soybean hybrid seeds have been developed, whereas they have not for wheat.

Table 2: Hypothesis Test Two.

| Crop-State | Number of Counties Rejecting Null | | |
| --- | --- | --- | --- |
| | Positive | Negative | Total |
| **Corn** | | | |
| Illinois | 0.0% | 16.5% | 16.5% |
| Indiana | 0.0% | 15.2% | 15.2% |
| Iowa | 0.0% | 5.1% | 5.1% |
| **Soybean** | | | |
| Illinois | 0.0% | 4.1% | 4.1% |
| Indiana | 2.4% | 2.4% | 4.9% |
| Iowa | 1.0% | 6.1% | 7.1% |
| **Winter Wheat** | | | |
| Kansas | 1.1% | 0.0% | 1.1% |
| Nebraska | 6.0% | 2.0% | 8.0% |
| Texas | 3.1% | 3.1% | 6.2% |

Note: Statistical significance evaluated at the 5% significance level using least Squares $t$-test with robust standard errors.

which uses the base case insurance valuation technique (simple trend), and (2) a private insurance company, which uses the proposed technique (two-trend). For the simulated game we follow Harri *et al.* (2011) and Ker and Coble (2003). If the two-trend estimates return statistically significant economic rents compared to the simple-trend estimates then we can conclude the two-trend leads to more efficient premium rate estimates.[10]

Note the simulation compares the ability of the empirical model to fit at the lower-end of the density tails and is not necessarily a reflection of the model to fit all of the data (whereas the in-sample likelihood ratio test conducted earlier was). Further, we can expect the simple trend method to have an *a priori* advantage in handling extremely low yield realizations (to see this, notice in Figures 2 and 3 how the lower tail of the two-trend model for 2010 has virtually no mass less than 100 bu./ac., while the simple trend model has going back to approximately 75 bu./ac.). This difference arises from the difference in heteroscedasticity treatments: the simple trend method treats heteroscedasticity directly, while the two-trend method has homoscedastic variances and accounts for heteroscedasticity indirectly through what we suspect is the underlying cause of heteroscedasticty: different technological trends in the component means. As a result, the two-trend method may consistently underprice the probability of an extremely low yield realization. That being said, we expect this difference in treatments to make only a marginal difference (the densities in this area are only 1 to 2%) and if two-trend method performs better than the simple trend despite this disadvantage, it provides further support to the case for a two-mean trend.

The design of the simulation imitates the decision rules of the Standard Reinsurance Agreement. Under the Standard Reinsurance Agreement, private insurance companies may effectively retain or cede insurance contracts of their choice *ex ante*.[11] Let $\hat{\pi}_{prv}$ be the estimated premium rate of the private insurance company and $\hat{\pi}_{gov}$ be the estimated premium

---

[10]The out-of-sample game is intuitively analogous to a game of chance: if I have superior knowledge of the true data generating process and therefore the outcome probabilities (more efficient premium rates), I should win more than my opponents over the long-run (economic rents).

[11]The SRA contains multiple funds and is more complicated but essentially a private insurance company can significantly reduce their exposure to unwanted policies.

rate of RMA. Define $\omega_{prv}$ as a discrete decision matrix for the private insurance company, where an element of $\omega_{prv}$ is equal to 1 for a retained contract and 0 for a ceded contract. The private insurance company will retain policies if and only if their rates are lower than the government rates: $\hat{\pi}_{prv} < \hat{\pi}_{gov}$. Let the $i^{th}$ row of $\omega_{prv}$ correspond to the $i^{th}$ year of the out-of-sample simulation and the $k^{th}$ column correspond to the county, then both $\omega$ matrices have dimensions $IxK$. The decision rule for any county $k$ and iteration $i$ is determined by the estimated rates of the two agents:

$$\omega_{prv} = \begin{cases} 0 & \text{if } \hat{\pi}_{prv} > \hat{\pi}_{gov} \\ 1 & \text{if } \hat{\pi}_{prv} < \hat{\pi}_{gov} \end{cases}$$

We run the simulation for 20 years with corn and soybeans but only ten years for wheat due to the limited length of the time series data available. Since Government-provisioned crop insurance in Canada is offered mainly at the farm-level, county-level results would be irrelevant and we exclude Ontario from the simulation. The $i^{th}$ iteration of the simulation corresponds to the time-period of the out-of-sample simulation. For example, consider the 1955-2011 corn and soybean data set with 20 iteration simulation, which implies $\{i_t\}_{t=1992}^{T=2011}$ where $t \in \mathbb{N}$. Considering any $n^{th}$ iteration $i_n$, the information set used to calculate premiums includes only yield values from $y_{n-2}$. In other words, the data set for a year $t$ and iteration $i_t$ uses years 1955 to $t - 2$.[12] All procedures and rates are re-estimated at each year. Loss ratios from the simulation for the two agents are calculated using actual yield realizations.

Table 3 summarizes the results of the out-of-sample simulation for all U.S. state-crop combinations at the 75% and 90% coverage level. Overall the private company loss ratio (using the two-trend method) is lower in 72.2% of state-crop combinations across both coverage levels. Of these, four state-crop combinations are statistically significant at the 5% significance level. Notably, all cases where the government (simple trend) loss ratio is lower the $p$-value is always far less than 0.900, indicating these results are not statistically signif-

---

[12]A two-step ahead forecast is used because the maximum amount of data available for an insuer to estimate the premium rate for year $n$ is $t - 2$.

Table 3: Out-of-sample simulation results for corn, soybeans and wheat.

| Set | Coverage Level | Retained by Private | Psuedo Loss Ratio Private | Psuedo Loss Ratio Government | p-value | % of Policies to Payout | Underpayment to program ($) | Overcharge to farmers ($) |
|---|---|---|---|---|---|---|---|---|
| **Corn** (20 years) | | | | | | | | |
| Illinois | 75% | 85.9% | 0.092 | 0.026 | 0.787 | 3.0% | -17,497,168 | 132,033 |
| | 90% | 87.8% | 0.287 | 0.465 | 0.012 | 18.6% | -24,797,891 | 873,866 |
| Indiana | 75% | 81.1% | 0.164 | 0.134 | 0.604 | 3.7% | -5,115,966 | 573,494 |
| | 90% | 82.4% | 0.395 | 0.434 | 0.373 | 19.8% | -8,628,025 | 763,927 |
| Iowa | 75% | 83.9% | 0.357 | 0.409 | 0.361 | 6.6% | -2,421,730 | 498,230 |
| | 90% | 86.8% | 0.406 | 0.444 | 0.291 | 12.2% | -3,697,009 | 523,887 |
| **Soybean** (20 years) | | | | | | | | |
| Illinois | 75% | 77.7% | 0.366 | 0.374 | 0.431 | 3.0% | -1,338,110 | 88,497 |
| | 90% | 67.4% | 0.540 | 0.474 | 0.695 | 12.9% | -2,318,225 | 170,473 |
| Indiana | 75% | 86.9% | 0.257 | 0.258 | 0.479 | 3.4% | -1,697,661 | 53,904 |
| | 90% | 78.4% | 0.611 | 0.746 | 0.201 | 19.6% | -3,826,867 | 125,585 |
| Iowa | 75% | 80.6% | 0.809 | 0.467 | 0.757 | 6.1% | -428,305 | 70,124 |
| | 90% | 78.7% | 0.751 | 0.911 | 0.076 | 16.1% | -755,864 | 84,005 |
| **Wheat** (10 years) | | | | | | | | |
| Kansas | 75% | 52.6% | 1.297 | 1.921 | 0.023 | 18.5% | -1,907 | 683 |
| | 90% | 41.9% | 1.268 | 1.399 | 0.217 | 33.0% | -3,729 | 1,382 |
| Nebraska | 75% | 40.8% | 0.309 | 0.747 | 0.070 | 5.3% | 0 | 5,125 |
| | 90% | 45.3% | 0.652 | 0.611 | 0.587 | 19.8% | 0 | 6,610 |
| Texas | 75% | 73.5% | 0.957 | 2.096 | 0.001 | 19.5% | -23,966 | 9,610 |
| | 90% | 69.1% | 1.112 | 1.593 | 0.001 | 43.8% | -46,529 | 18,851 |

Note: Out-of-sample simulation for 20 years with corn and soybeans (where county-level data series are 57 years long) but only 10 years for wheat due to limited data availability (county-level data series are 41 years long). Counties are equally-weighted. Overcharge and underpayment are low for wheat due to very little participation in area-yield insurance programs.

icant.[13] The two-trend method also appears to typically perform better at higher coverage levels and when the per cent of policies paying out is higher. The purpose of the out-of-sample simulation was to illustrate the economic applications of the proposed empirical model. Using crop insurance premium rates as an example, the results of the out-of-sample simulation indicate the proposed empirical model can lead to economically and statistically significant improvements in crop insurance premium rates.

## Conclusions

This manuscript presented a new approach to considering technological change in crop yields that we would suggest is more consistent with how research in plant science is undertaken; the development of new technologies is generally targeted at subpopulations of the yield distribution rather than a uniform upward shift in the yield distribution. Examples include the development of *racehorse* seeds and drought resistant seeds. We proposed a mixture model that explicity accounts for distinct subpopulations in the yield distribution.

The proposed model is sufficiently flexible to accomodate a number of different distributional shapes that the literature has suggested for modelling crop yields. The proposed model also allows for differing rates of technological change in differing subpopulations of the yield distribution which opens up a number of interesting hypotheses. Our results provide compelling evidence to suggest technological advancement has increased greater in *regular* versus *poor* years. A consequence of differing rates of technological change is heterosedasticity which is prevalent in yield data. The test results also suggest the predominant approach in the literature -- time-conditioning yields at the mean -- may not fully capture the effect of technological change on yields. These test results are also supported by the results of our simulated game of rating insurance contracts whereby economically and statistically significant rents can be gained by using the two-trend normal mixture model to averse select against the simple-trend normal mixture model. The results of hypothesis two are also in-

---

[13]$p$-value $\rightarrow 1$ indicate a randomized allocation of contracts would always do better than the two-trend approach.

teresting and suggest that despite climate change concerns, we do not find evidence that the probability of a *regular* or *poor* year is changing through time.

One caveat regarding the yield data is worth repeating; producers may choose only a given set of technologies per growing season and as such the realized yield data is only representative of technologies adopted and not the full set of technologies available to producers. Given the development of *racehorse* seeds and the widespread availability of crop insurance, it is highly unlikely that producers would adopt technologies that would perform adequately in sub-optimal conditions at the expense of signficantly lower yields in optimal conditions. As such, the realized yield data is only a function of **adopted** technologies and by default so too are the results presented here.

# References

Atwood, J., S. Shaik, and M. Watts. 2003. "Are crop yields normally distributed? A reexamination." *American Journal of Agricultural Economics* 85 (4):888 – 901.

Bilmes, J.A. 1998. "A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models." Working paper No. TR-97-021, International Computer Science Institute: Berkeley, CA.

Cabas, J., A. Weersink, and E. Olale. 2010. "Crop yield response to economic, site and climatic variables." *Climatic Change* 101(3):599–626.

Claassen, R., and R.E. Just. 2011. "Heterogeneity and distributional form of farm-level yields." *American Journal of Agricultural Economics* 93 (1):144 – 160.

Coble, K.H., R. Heifner, and M. Zuniga. 2000. "Implications for crop yield and revenue insurnace for produces hedging." *Journal of Agricultural and Resource Economics* 23:539–551.

Day, R.H. 1965. "Probability distributions of field crop yields." *Journal of Farm Economics* 47 (3):713 – 741.

Du, X., D.A. Hennessy, and C.L. Yu. 2012. "Testing Day's conjecture that more nitrogen decreases crop yield skewness." *American Journal of Agricultural Economics* 94(1):225–237.

Featherstone, A.M., and T.L. Kastens. 2000. "Non-parametric and semi-parametric techniques for modeling and simulating correlated, non-normal price and yield distributions: Applications to risk analysis in Kansas agriculture." *Journal of Agricultural and Applied Economics* 32 (2):267 – 281.

Finger, R. 2010. "Revisiting the evaluation of robust regression techniques for crop yield data detrending." *American Journal of Agricultural Economics* 92 (1):205 – 211.

Gallagher, P. 1987. "U.S. soybean yields: Estimation and forecasting with nonsymmetric disturbances." *American Journal of Agricultural Economics* 69 (4):796 – 803.

Goodwin, B.K., and A.P. Ker. 1998. "Nonparametric estimation of crop yield distributions: Implications for rating group-risk crop insurance contracts." *American Journal of Agricultural Economics* 80 (1):139 – 153.

Halcrow, H.G. 1949. "Actuarial structures for crop insurance." *Journal of Farm Economics* 31 (3):418 – 443.

Harri, A., K. Coble, A.P. Ker, and B.J. Goodwin. 2011. "Relaxing heteroscedasticity assumptions in area-yield crop insurance rating." *American Journal of Agricultural Economics* 93 (3):707 – 717.

Harri, A., C. Erdem, K.H. Coble, and T.O. Knight. 2008. "Crop yield distributions: A reconciliation of previous research and statistical tests for normality." *Review of Agricultural Economics* 31 (1):163 – 182.

Just, R.E., and Q. Weninger. 1999. "Are crop yields normally distributed?" *American Journal of Agricultural Economics* 81 (2):287 – 304.

Ke, W., and Z. Qiao. 2012. "Influence of flexible crop yield distributions on crop inurance premium rate - A case study on cotton insurance in three areas of Xinjiang." *Asian Agricultural Research* 4(5):7–12.

Ker, A.P., and K. Coble. 2003. "Modeling conditional yield densities." *American Journal of Agricultural Economics* 85 (2):291 – 304.

Ker, A.P., and P. McGowan. 2000. "Weather-based adverse selection and the U.S. crop insurance program: The private insurance company perspective." *Journal of Agricultural and Resource Economics* 25(2):386–410.

Knight, F.H. 1921. *Risk, Uncertainty and Profit*, New York: Reprints of Economics Classics, chap. Structures and methods for meeting uncertainty. pp. 233 – 263.

Leisch, F. 2003. "FlexMix: A general framework for finite mixture models and latent class regression in R." Report Series No. 86, Adaptive Information Systems and Modelling in Economics and Management Science: Wien, Austria.

Lobell, D.B., K.G. Cassman, and C.B. Field. 2009. "Crop yield gaps: their importance, magnitudes and causes." Working paper, NCESR Publications and Research. Paper 3.

Luttrell, C.B., and R.A. Gilbert. 1976. "Crop yields: Random, cyclical or bunchy?" *American Journal of Agricultural Economics* 58 (3):521 – 531.

Miller, E.M. 1977. "Risk, uncertainty and divergence of opinion." *The Journal of Finance* 32 (4):1151 – 1168.

Miranda, M.J. 1991. "Area-yield crop insurance reconsidered." *American Journal of Agricultural Economics* 73 (2):233–242.

Moss, C.B., and J.S. Shonkwiler. 1993. "Estimating yield distributions with a stochastic trend and non-normal errors." *American Journal of Agricultural Economics* 75 (5):1056 – 1062.

Nelson, C.H. 1990. "The influence of distributional assumptions on the calculation of crop insurance premia." *North Central Journal of Agricultural Economics* 12(1):71–78.

Nelson, C.H., and P.V. Preckel. 1989. "The conditional beta distribution as a stochastic production function." *American Journal of Agricultural Economics* 71:370 – 378.

Norwood, B., M.C. Roberts, and J.L. Lusk. 2004. "Ranking crop yield models using out-of-sample likelihood functions." *American Journal of Agricultural Economics* 86 (4):1032 – 1043.

Park, R.E. 1966. "Estimation with Heteroscedastic Error Terms." *Econometrica* 34(4):888.

Pease, J.W. 1992. "A comparison of subjective and historical crop yield probability distributions." *Southern Journal of Agricultural Economics* Dec.:23 – 32.

Ramirez, O.A. 1997. "Parametric model for simulating heteroskedastic, correlated, nonnormal random variables: The case of Corn Belt corn, soybean and wheat yields." *American Journal of Agricultural Economics* 79 (1):191 – 205.

Ramirez, O.A., and T. McDonald. 2006. "Ranking crop yield models: A comment." *AJAE* 88 (4):1105 – 1110.

Ramirez, O.A., S. Misra, and J. Field. 2003. "Crop-yield distributions revisited." *American Journal of Agricultural Economics* 85 (1):108 – 120.

Schmitz, A., C.B. Moss, T.G. Schmitz, H.W. Furtan, and H.C. Schmitz. 2010. *Agricultural Policy, Agribusiness, and Rent-Seeking Behaviour (2nd ed.)..* Toronto, ON: University of Toronto Press.

Sherrick, B.J., F.C. Zanini, G.D. Schnitkey, and S.H. Irwin. 2004. "Crop insurance valuation under alternative yield distributions." *American Journal of Agricultural Economics* 86 (2):406 – 419.

Skees, J.R., J.R. Black, and B.J. Barnett. 1997. "Designing and rating an area yield crop insurance contract." *American Journal of Agricultural Economics* 79 (2):430 – 438.

Skees, J.R., and M.R. Reed. 1986. "Rate making for farm-level crop insurance: Implications for adverse selection." *American Journal of Agricultural Economics* 68 (3):653 – 659.

Swinton, S.M., and R.P. King. 1991. "Evaluating robust regression techniques for detrending crop yield data with nonnormal errors." *American Journal of Agricultural Economics* 73 (2):446 – 451.

Tack, J., A. Harri, and K. Coble. 2012. "More than mean effects: Modeling the effect of climate on the higher order moments of crop yields." *American Journal of Agricultural Economics* 94(5):1037–1054.

Turvey, C., and J. Zhao. 1999. "Parametric and non-parametric crop yield distributions and their effects on all-risk crop insurance programs." Unpublished, University of Guelph Working Paper.

Woodard, J.D., and B.J. Sherrick. 2011. "Estimation of mixture models using cross-validation optimization." *American Journal of Agricultural Economics* 93 (4):968 – 982.